LA-UR-20-24084
Approved for public release; distribution is unlimited.

Title: Hydrodynamic and Radiographic Toolbox (HART)

Author(s): Klasky, Marc Louis; Nadiga, Balasubramanya T.; Disterhaupt, Jennifer Lynn Schei; Guardincerri, Elena; Carroll, James Lamond; Pfister, Luke Anthony; Hovey, Luke Douglas; Skau, Erik West; Salazar, Gary P.; Warthen, Barry J.; Mendez, Albert Jacob; Schulze, Martin E.; Coleman, Joshua Eugene; Rose, Chris Randall; Espy, Michelle A.; Wohlberg, Brendt Egon; Han, Yoseob; Wilcox, Trevor; Mockler, Theodore; Hungerford, Aimee L.; Ravishankar, Saiprad; et al.

Intended for: Collaboration with Universities

Issued: 2020-06-04

Hydrodynamic and Radiographic Toolbox (HART)

PI: Marc Klasky, P-23
Co-PI: Balu Nadiga, CCS-2
Team Internal: Jennifer Schei Disterhaupt, P-23; Luke Pfister, P-23;
Elena; Guardincerri, P-23; James Carroll, P-23; Luke Hovey, P-23
Erik Skau CCS-3;
Gary Salazar, J-4; Barry Warthen, J-4; Jacob Mendez, J-4;
Marin Schulze J-5, Josh Coleman J-5, Chris Rose J-5
Michelle Espy, E-6;
Brendt Wohlberg, T-5; Yoseob Han, T-5;
Aimee Hugerford, XCP-DO;
Trevor Wilcox, XTD-SS; Ted Mockler, XTD-SS;
Team External: Saiprasad Ravishankar, MSU; Michael McCann, MSU, Yoram Bresler,
UIUC; Berk Iskender, UIUC; Anish Lahiri, UM; and Charlie Bouman, Purdue

May 2020

Abstract

With a multi-lab and university team, we propose to develop new methods for a Hydrodynamic and Radiographic Toolbox (HART) that will enable a fuller and more extensive use of experimental radiographic data towards better characterizing and reducing uncertainties in predictive modeling of weapons performance. We will achieve this by leveraging recent developments in areas of computational imaging, statistical and machine learning, and reduced order modeling of hydrodynamics. In terms of software practices, by partnering with XCP (RISTRA Project) we will conform to recent XCP standards that are inline with modern software practices and standards and ensure compatibility and ease of inter-operability with existing codes.

Three new activities under this proposal include (a) the development and use of deep learning-based surrogates to accelerate reconstruction and variational inference of density fields from radiographs of hydrotests, (b) a model-data fusion strategy that couples deep learning-based density reconstructions with fast hydrodynamics simulators to better constrain the reconstruction, and (c) a method for treating asymmetries using techniques adopted from limited view tomography.

All three new activities will be based on improved treatment of scatter, noise, beam spot movement, detector blur, and flat fielding in the forward model, and a use of sophisticated priors to aid in the re construction. The improvements to the forward model and improved algorithmic design of the reconstruction when complete will be contained in the iterative reconstruction code SHIVA—a code project that we have recently initiated. The many ways in which machine learning can be used in the reconstruction work will be contained in a code HERMES that has been initiated with DTRA support. For example, the significant levels of acceleration that will likely be achieved by the use of machine learning techniques will permit us (and are required) to quantify uncertainties in density retrievals. Next, the two-way coupling between density reconstruction and model-based simulation of the hydrodynamics will be contained in code EREBUS, and will permit a fuller realization of the potential of the data to constrain the hydrodynamic model and better address issues related to asymmetries in the problem. Finally, we anticipate that the better consistency with physics achieved in our reconstructions will allow them to be used by X-Division more so than today.

1 Problem Statement

The Dual Axis Radiographic Hydrodynamic Test Facility (DARHT) plays a key "last stop" role in the National Nuclear Security Administration's approach to Stockpile Stewardship in the absence of nuclear tests [14]. However,

the noisy and complex multi-scale and multi-physics environment of dynamic tests necessitates careful, quantitative analysis of the radiographic data obtained from the facility to accurately reconstruct the density field and obtain the various quantities of interest, such as the location of the metal edge, location of the shock, and other details of the density distribution, along with three-dimensionality.

2 Program Needs and Gaps

The current approach to density reconstruction from radiographic data was conceived of in the early 90s and the attendant "Bayesian Inference Engine" (BIE) software has been in use for over two decades [41, 42]. Having performed over 90% of the radiographic analyses for DARHT during the past twenty years, we have identified a number of shortcomings with the current approach that prevent it from making full use of the radiographic data. These shortcomings range from slowness of the analysis (it takes months to analyze a hydrotest), reliance on axisymmetry [51, 53, 98], lack of a physics model to describe scatter [98], lack of incorporation of an underlying model linking the temporal dynamics [68], no formal treatment of noise [98], as well as largely empirical models to describe the detector blur [98]. These shortcomings limit the ability to provide robust estimates of uncertainty [19–21, 98, 99]. Furthermore, because of its use of antiquated and proprietary software tools, developments in related fields over the intervening couple of decades that can help overcome these shortcomings cannot be readily brought into the current approach/framework (BIE).

3 Research Approach

In inverse problems such as density reconstructions from radiography data, it is typical to build an "outer loop" over a physics-based forward model to ensure physical consistency of the found solution. In a first view of the density reconstruction problem, e.g., as with a single radiographic image, the forward model may be thought of as comprising only the radiographic measurement system. However, when we have a temporal sequence of radiographic data, it is possible, and perhaps necessary, to enlarge the notion of the forward model to include the hydrodynamic system in addition to the radiographic measurement system. Finally, when a temporal sequence of images along one axis is augmented with an image along a second axis, the forward model has to be similarly augmented.

This hierarchical view of DARHT radiographic data forms the basis for our three proposed activities: (a) the development and use of deep learning-based surrogates to accelerate accurate reconstruction and variational inference of density fields from radiography of hydrotests, (b) a model-data fusion strategy that couples deep learning/and or Koopman Operator projections to allow for density reconstructions with fast hydrodynamics simulators to better constrain the reconstruction, and (c) a method for treating asymmetries using techniques adopted from limited view tomography. All three activities will be based on improved treatment of scatter, noise, beam spot movement, detector blur, and flat fielding in the forward model, and a use of sophisticated priors to aid in the reconstruction.

The first proposed activity, in addition to providing all current capabilities, will likely perform better and permit a characterization of uncertainties in density retrievals. The two-way coupling between density reconstruction and model-based simulation of the hydrodynamics in the other activities will permit a fuller realization of the potential of the data to constrain the hydrodynamic model and better address issues related to asymmetries in the problem. Finally, we anticipate that the better consistency with physics achieved in our reconstructions will allow them to be used by X-Division more so than today.

These activities will be integrated in a Hydrodynamic and Radiographic Toolbox (HART). A schematic of this toolbox is shown in Figure 1. We present these approaches in increasingly levels of complexity to address both the direct replacement of BIE with a modern fully parallel iterative reconstruction algorithm (SHIVA) including enhancements in virtually all aspects of the physics including a complete descattering algorithm, as well as enhancements in optimization routines, and the use of plug-and-play priors as described below. The code HERMES addresses our desire to leverage our successful neural networks in performing both descattering as well as density reconstructions in real time. EREBUS expands upon the HERMES code by coupling a hydrodynamic prior to allow for the most accurate density fields to be obtained from the data in a manner consistent with an underlying hydrodynamic path.

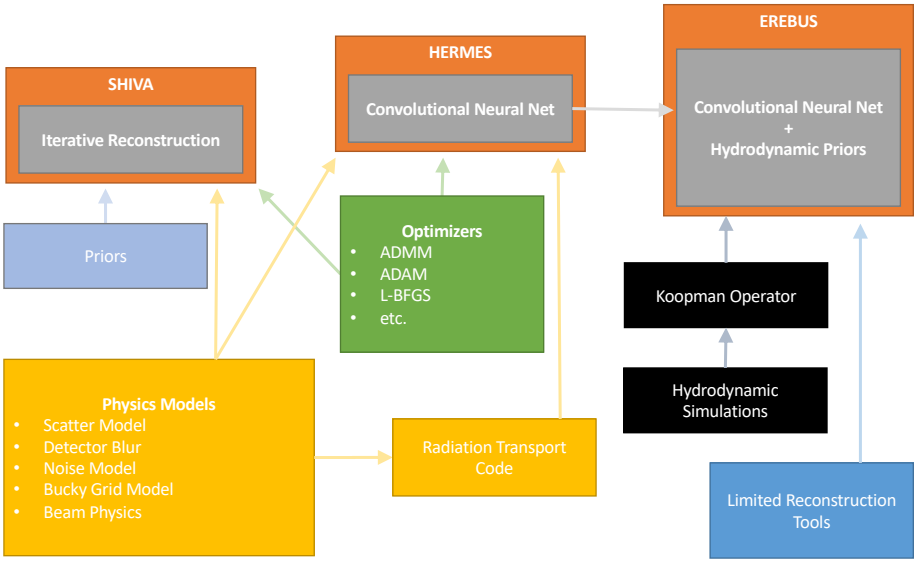

Figure 1: A schematic of the toolbox, which comprises the codes Shiva, Hermes, and Erebus.

3.1 Improvements to the Forward Model

Density reconstructions obtained from radiographs of dynamic tests are complicated by the noisy and complex multi-scale and multi-physics environment. As such, while there are robust features in such reconstructions such as the shock location and edge location that may be extracted with quantifiable uncertainty from a combination of static experimental data as well as from sensitivity analysis, other aspects such as details of the density field are subject to greater uncertainty. The latter uncertainties arise from our inability to exactly represent various aspects of the radiographic measurement system, such as scatter, beam spot movement, beam-target interactions, beam dynamics repeatability, and aspects of the image formation process in the forward model.

The most crucial part of any model, and consequently the radiographic inversion, is the understanding of the underlying physics. As previously mentioned, the BIE forward model almost exclusively relies upon empirical models and thereby conflates physics. Consequently, this code, in its current state, cannot be relied upon to perform density reconstructions to the level of accuracy needed to inform the hydrodynamic codes. Accordingly, we have developed physics models to replace many of the missing or incorrect models utilized in the previous radiographic reconstruction approach. These physics models will be used in the iterative reconstruction approach. In addition, they will be applied to the data generation for the machine learning approach.

We now discuss our current physics models that we propose to implement into the forward model to enhance the data fidelity term and improve density reconstructions.

Detector Physics and Noise

Detector blur consists of the blurring of X-rays and that of optical light photons, both within the scintillator, and the blur due to the lens system coupling the scintillator to the camera. Having performed transport simulations of scintillator X-ray blur using MCNP6 [38], and that of optical light photons using GEANT4 [2], and having measured blur due to the lens, we will combine models for each of these components into an overall model for detector blur. A similar procedure of detailed simulation and measurement will be used to develop and calibrate a noise model.

Beam Physics

The inference of the DARHT source spot is currently performed using the BIE using resolution resolution targets fielded several days prior to the dynamic experiment [28, 45]. However, in so far as the source may both move and also possess different characteristics during the dynamic experiment than those inferred during the characterization shots performed prior to the dynamic experiment we propose to re-examine this approach. The approach will incorporate both measurements obtained from the beam focusing diagnostics during the dynamic experiment as well as beam measurements, pinhole measurements leading up to the shot. In addition, we propose to perform simulations using Particle in Cell (PIC) codes [26, 89], MCNP [38], GEANT [2], as well as LASNEX [112] to model the beam characteristics. Using these calculations, we propose to perform sensitivity studies to examine the impact of variations in physical machine parameters on the beam characteristics *i.e.*, shape, size, energy, current, and location. This will allow for the establishment of actual requirements on the physical parameters that control the beam spot formation and thereby limit potential variations during the dynamic experiment. Indeed, by establishing these requirements we will reduce the variability of the beam spot characteristics and reduce the need for the empirical model utilized within the current radiographic model. Presently, the model of the gamma spectrum is assumed to be the same for all four pulses on DARHT Axis II. Simulations [92] have shown that there can be significant variations (20%) in the gamma ray spectrum on DARHT Axis II for the four pulses and that these variations are sensitive to the dose format and the target geometry. We plan to incorporate these variations in the energy spectra into our radiographic models. Furthermore, we have proposed using the Compton Spectrometer as a means to experimentally validate the energy spectra. Finally, we will address the alignment between the object and source as well as the source to the GRC. If successful, this could potentially remove the angular dependence of the beam striking the target.

Flat Field Treatment

A flat field is an image taken without an object in the line-of-sight but with a significant amount of attenuation. Ideally, this would result in an image with no transverse structure. The measured image shows the impurities/variations of the scintillator pixels, imperfections in the septa between the scintillating material and the camera, and, more importantly, the gain variation of each camera pixel. In current analysis techniques, the dynamic image is divided by the flat field image to remove the high frequency intensity modulation and the septa in the dynamic image. There are several drawbacks to this method. First, the flat field image is typically taken 1–3 days prior to a dynamic experiment due to physical constraints. As discussed previously, the beam characteristics *e.g.*, spot size and location may change between the flat field and dynamic images, which introduces errors in the preprocessed image. Second, the flat field image contains a non-stationary scatter field. It should be noted that in the BIE forward model, scatter in the flat-field was incorrectly assumed to be zero [65]. This phenomenon has been unaccounted for and has caused problems in past density reconstructions [66]. Finally, dividing a noisy dynamic image by a noisy flat field image increases the both the magnitude and complexity of the noise in the final image, making feature and density identification more difficult.

We propose to investigate more sophisticated methods of preprocessing the dynamic image instead of dividing by the flat field image. The dynamic image can be converted into frequency space, and the scintillator septa can be removed by clipping that particular frequency of the image. The patchiness of the scintillator can be learned using machine learning techniques and applied to the dynamic image using feature location identification.

Scatter Closure Relationship

We will improve upon the current practice of modeling scatter as an additive (polynomial) field independent of the density field [98] by developing a closure-based model that relates the scattered radiation to the density field being radiographed [60, 67, 69]. This development will be based on detailed computations using MCNP6 and will leverage our NNSA/DTRA-funded efforts to descatter radiographic images [65, 69] (briefly described below).

Briefly, our kernel based scatter models are based on a general representation given by,

$$S = f(D), \tag{1}$$

where S is scatter and D is the direct transmission in the absence of scatter. Furthermore, we have developed models that represent the functional by,

$$S = (C \odot D) * K_t(D), \tag{2}$$

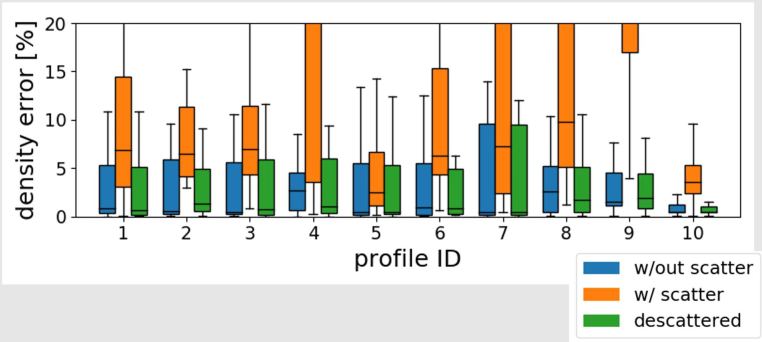

Figure 2: Scatter removal using the kernal-based physics-based model reduces density reconstruction errors, as illustrated on 10 separate objects.

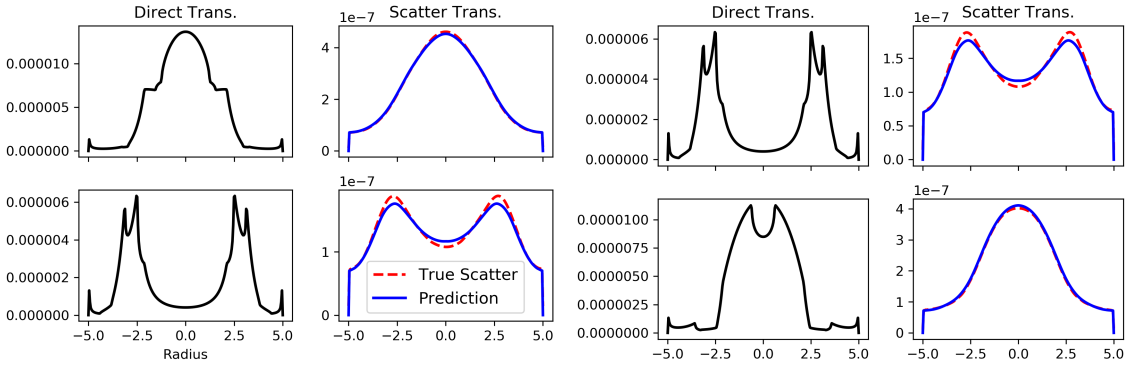

Figure 3: A Convolutional Neural Network (CNN) was trained to learn the mapping from direct radiation to scattered radiation. The black traces show the direct signal, the blue traces show the true scattered signal, and the red dashed traces show the predicted scatter signal from the CNN.

where $\odot$ denotes pixel-wise multiplication and $*$ denotes convolution. Here, C is a multiplicative field that weights different rays in the transmission according to the areal mass of material along the ray so that a larger areal mass traversed would mean more scattered photons, and K is a scatter kernel (estimated from MCNP scatter simulations of axis-symmetric objects). Results using this approach are shown in Figure 2.

More sophisticated scatter models assign different scatter kernels for different ranges of areal masses. We have explored models of the form,

$$S = S_{\text{coherent}}(D) + \alpha(D) S_{\text{compton}}(D), \tag{3}$$

where S_{coherent} and S_{compton} are the estimates of coherent and first-order Compton scatter modeled separately (*e.g.*, via local convolutional kernels and/or with multiplicative fields), and $\alpha(D)$ is an additional scaling of the Compton scatter to account for higher-order Compton scatters (*e.g.*, coherent scattered photons that undergo subsequent Compton scattering, etc.).

Finally, Convolutional Neural Network (CNN)-based scatter models learned from data sets of MCNP scatter simulations have also been examined, where the function f in (1) is a neural network with learnable weights [60, 69]. An example of these results is shown in Figure 3. Future research is anticipated in this area to enable a CNN architecture that is based on the physics of scattering processes, similar to networks parameterized by the Neumann expansion for use in electromagnetic scattering [36]. These models may have the natural advantage of being more generalizable and easier to train than general CNNs.

The above scatter models will be used to descatter experimental radiographs and perform density reconstructions. Both fixed point and optimization approaches will be used for descattering. The fixed point approach models the transmission as,

$$T = D + f(D), \tag{4}$$

where T is the total transmission, D is the direct transmission and $f(D)$ is a scatter model. This method starts with an initial estimate,

$$D \approx D^0 = T, \tag{5}$$

and iterates between computing scatter at the k-th iteration as,

$$S^k = f(D^{k-1}), \tag{6}$$

and estimating the direct transmission as,

$$D^k = T - S^k = T - f(D^{k-1}). \tag{7}$$

An additional projection of D^k onto the non-negative orthant may be needed to avoid unrealistic estimates. This fixed point algorithm is fairly simple, computationally very efficient, and is run until convergence (*i.e.*, until the error between subsequent estimates of D are sufficiently small). However, its robustness in the presence of detector noise is not well studied.

An alternative approach we will investigate for descattering is solving the optimization problem,

$$\hat{D} = \underset{D \geq 0}{\arg\min} \| T - D - f(D) \|^2 + \lambda R(D), \tag{8}$$

where we aim to find the non-negative direct signal D that provides the best fit (in the l_2 sense) of,

$$D + S = D + f(D), \tag{9}$$

to the transmission T. An additional regularizer, or prior, on D is included as is common in optimization-based reconstruction. The squared Euclidean loss can be generalized to other loss functions, depending on the noise model for measurements, *e.g.*, Poisson noise.

The scatter models along with the descattering algorithms may be implemented into both our iterative reconstruction algorithm as a prior as well as in our Convolutional Neural Network Reconstruction algorithm to correct one of the fundamental flaws in the current radiographic approach.

Using the scatter models and accompanying descattering algorithm, we have obtained very promising results that demonstrate the accuracy of scatter modeling as well as descattering techniques (Figures 2).

Finally, it should also be noted that one of the convolutional kernel-based descattering methods was previously utilized on a recent hydro-test to allow for the best density reconstruction with Axis 2 radiographic data to date [58, 63].

Bucky Grid Model

An associated aspect of modeling scatter is the treatment of the Axis 1 Bucky Grid. The current empirical polynomial model of the Bucky Grid, when compared to a physics-based scattering model, has been shown to result in errors in the density reconstructions of the French Test Object that are actually larger than those without the Bucky Grid [56,57]. As such, we will replace the empirical model with models that we will develop using kernel-based and machine learning approaches after conducting detailed simulations.

Having addressed the major improvement to the physics that will be utilized in all of our algorithms, we now present details of the three respective codes that comprise HART (Figure 1). These approaches in increasingly levels of complexity include the direct replacement of the iterative forward model of the BIE with a modern, fully parallel, iterative reconstruction algorithm (SHIVA). It will contain enhancements in virtually all aspects of the physics including a complete descattering algorithm, enhancements in optimization routines, and the use of plug-and-play priors as described below. The code, HERMES, addresses our desire to leverage our successful neural networks in performing both descattering as well as density reconstructions in real time. EREBUS expands upon the HERMES code by coupling a hydrodynamic prior to allow for the most accurate density fields to be obtained from the data in a manner consistent with an underlying hydrodynamic path.

3.2 Improvements to the Outer Loop

The traditional approach to solving the radiographic inversion problem with a complex non-linear/non-convex forward model is to use iterative reconstruction techniques. Mathematically we may express this problem as:

$$\hat{x} = \arg\min_x F(x) + \sum_i \alpha_i R_i(x) \tag{10}$$

where $F(x)$ is the data fidelity term capturing the forward model of the imaging process and the statistical models of measurements and noise, and the $R_i(x)$ are regularizer models for the desired image. The scalar parameters α_i control the relative strength of the respective regularizers.

The determination of the correct density distribution from (10) is contingent not only upon formulating a proper physics-based forward model, but also imposing appropriate regularization and effective schemes to solve the optimization problem.

Priors and Regularized Inversion

Regularization, *i.e.*, the use of a prior, is vital when solving ill-posed inverse problems, such as attempting to perform radiographic inversions. The regularizer mitigates the ill-posedness of the inverse problem by encouraging solutions to obey a particular model and/or have certain desirable properties [6]. Expressive prior models that properly capture the salience of the underlying data lead to improved solution fidelity.

In the over twenty years since Hansen formulated the BIE [41, 42], significant progress has been made in developing powerful new priors for imaging applications by both our team and the larger community. Indeed our team members have been significant contributors in this area of research.

Many popular prior models are based on sparse representation. For example, smoothness can be viewed as sparsity of the gradient under the finite-difference operation [32, 79]. More sophisticated priors are based on sparsity under directional multi-scale decompositions such as wavelets, shearlets, and curvelets [16, 17, 34, 73, 79, 102]. More recently, robust prior models have been learned in conjunction with the image reconstruction process by several team members [23, 44, 87, 88, 100, 109, 110].

In parallel, there have been significant advances in the solution of the simplest inverse problem: image denoising. These denoising algorithms typically combine sparse representations with non-local self-similarity [13, 25]. More recently, denoising algorithms based on Deep Convolutional Neural Networks have obtained performance beyond what was previously thought possible [111]. However, as these denoising algorithms typically incorporate image models in an implicit manner, it is not obvious how to leverage them to solve a general inverse problem of the form (10); that is, the denoising algorithm does not correspond to a particular regularization function $R(x)$.

Plug-and-Play Priors is a recent technique that exploits ideas from proximal optimization algorithms to combine a solver for the forward model with a generic image denoising algorithm in a modular way [104]. Using the Plug-and-Play Prior approach, powerful image denoising algorithms can be utilized to solve a generic inverse problem of the form (10).

The Plug-and-Play framework [95, 97, 104] is based on a direct application of the Alternating Directions Method of Multipliers (ADMM) [105] that has recently become popular for the solution of a variety of MAP estimation/ regularized inverse problems. The ADMM works by first splitting the state variable so as to decouple the prior and forward model terms of MAP estimation problem. The application of the ADMM technique to the resulting constrained minimization problem then results in two decoupled optimizations, one for the forward model and one for the prior model. We note that this allows for a completely decoupled software implementation.

Using these advanced techniques we submit that we can significantly improve upon the current limited techniques within the existing radiographic code.

Optimization

Inverse problems and machine learning problems typically involve minimizing either convex or nonconvex cost functions [32], sometimes with additional constraints (e.g., non-negativity of material density in radiographic imaging, constraints based on hydrodynamics laws, etc.). Before proceeding it should be noted that if the objective is not convex, only a local minimizer can be found. Gradient-based algorithms (e.g., gradient descent, conjugate gradients) or quasi-newton methods (e.g., BFGS [35], L-BFGS [15]) are often exploited for unconstrained

optimization involving convex, smooth objectives. We will also explore the use of Nesterov's proposed fast gradient method (FGM) [84]. Moverever, algorithms like the optimized gradient method (OGM) [50], which guarantees a 2× worst-case speedup over the FGM, will also be examined.

We will also examine majorize-minimize (MM) and proximal mapping-based methods that enable alternatives such as the fast iterative shrinkage-thresholding algorithm (FISTA) [5] and the proximal optimized gradient method [50], both of which have been shown to provide convergence of the cost. Variable-splitting based methods such as the the augmented Lagrangian method [1] and the alternating direction method of multipliers (ADMM) [12] have also been highly popular for convex and nonsmooth or constrained optimization problems. The MM methods encompass a broad array of algorithms that are helpful even in cases when the objective is non-convex. Quadratic surrogate algorithms is one class of such methods that rely upon optimization transfer to solve non-convex optimization problems.

Finally, distributed and asynchronous optimization algorithms [4, 86, 96] are garnering increasing attention recently and could be applied to large-scale radiographic inverse problems and machine learning tasks. Additionally, we will make use of automatic differentiation techniques so that new cost functions and regularizers can be easily implemented without the need to manually compute gradients [8, 43, 49].

Machine Learning based Reconstruction

While the traditional iterative reconstruction based on a forward model will be a part of the toolbox, we additionally propose to use machine learning approaches to accelerate the reconstruction process. As shown below, our preliminary results are encouraging and demonstrate the feasibility of the approach. We now briefly summarize our approach.

Recent advances in machine learning, in particular deep neural networks (DNNs), have translated into state of the art performance in a wide variety of computational imaging tasks (e.g., see the recent review by Gilton *et al.* [36]). These methods vary in how knowledge of the image formation process (i.e. the forward model and noise distribution) is incorporated into the learning process. One approach is to learn a DNN to implement the reconstruction mapping. This approach relies entirely on the learning algorithm to infer all relevant information from given training pairs, and does not explicitly rely on the forward model or noise distribution. The application of such a DNN is orders of magnitude faster than traditional iterative reconstruction techniques [18].

An alternative approach is to incorporate knowledge of the forward operator and physics directly in the structure of the neural network. One way to accomplish this is by first applying an approximate inverse of the forward model to measurements, then passing the result into a neural network [47]. An alternative approach utilizes neural networks in the Plug-and-Play framework.

In addition, we plan to explore the use of physics-based machine learning approaches by incorporating the projector into the machine learning as well as more fundamentally-based architectures motivated by techniques comparable to the method of moments [30]. These techniques will provide an underlying tie to physics, which should enable learning from comparatively few training pairs.

We will leverage our work performed for DTRA/NNSA which has achieved tremendous progress in performing reconstructions using a pure Convolutional Neural Network (CNN) to learn direct-to-density and direct-to-scatter mappings for Emergency Response Applications. We have combined these mappings with a fixed-point optimization scheme to develop a state-of-the-art descattering algorithm.

We now present some results from our DTRA/NNSA work to better illustrate the approach [60]. We have constructed test problems in which five materials out of 18 available materials were selected at random to build a series of five concentric spheres with an outer radius ranging from 10- 50 cm. Each shell thickness was also randomly selected. Simulated radiographs of direct and scattered radiation were then generated using MCNP6 [38] with a poly-energetic bremsstrahlung source to both train the networks as well as test the ability to reconstruct the density objects. The reconstructed radiographs were also performed in the presence of noise, departures from assumed energy spectra, and blur. The histogram of RMS density reconstruction errors, along with representative samples, is shown in Figure 4. Note that the after training the neural network, the results are obtained almost instantaneously.

We have also constructed a neural network to learn the scattered radiation from the direct radiation which may be utilized in the descattering algorithms described above. Results from this CNN are shown in Figure 3.

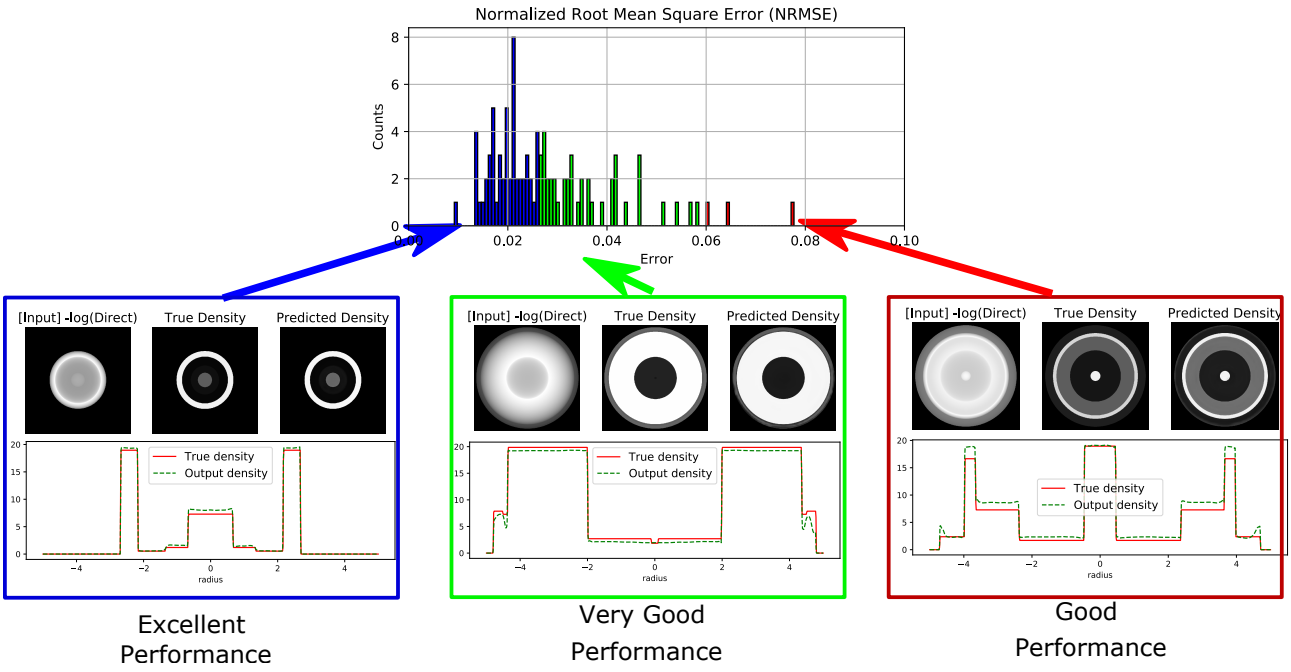

Figure 4: Density reconstruction from polyenergetic direct measurements using CNNs yields promising results. Top: Histogram of normalized RMS density reconstruction errors for 100 test objects. Middle: The log of direct radiation is the input to the CNN. We show the true density and the predicted density fields. Bottom: Radial line-out of the true (red) and predicted (green) density distributions.

3.3 Dynamic Reconstruction: Coupled Radiographic and Hydrodynamic Inversion

While the density reconstructions obtained using our CNNs look very promising to reconstruct single radiographic slices from simulations, our extensive experience in reconstructing experimental radiographic data suggest that the inability to perfectly descatter as well as a multitude of other experimental artifacts *i.e.*, beam spot movement, uncertainties in both accelerator behavior etc., will likely introduce error into our density reconstructions. Furthermore, our experience also has proven that inclusion of "fitting parameters" utilized within the BIE methodology does not afford either a physics or mathematically justifiable model that enables the determination of density fields with sufficient accuracy to inform the hydrodynamic codes and satisfy X-Divison.

Given this complicated scenario, we think it is advantageous to consider separate methodologies in an initial phase that on the one hand builds on robust aspects of the reconstruction and on the other hand develops other procedures to deal with and use the more uncertain aspects of the reconstruction. For example, with such an approach, robust aspects of the reconstruction can be leveraged to bracket certain parametric uncertainties in the hydro model. Indeed, such an approach based on realizing consistency between the radiographically diagnosed density field and its hydrodynamically simulated counterpart requires us to not only use fast methods for density reconstruction that leverage deep learning methodologies [37, 71], but also as importantly utilize reduced representations of the hydrodynamics using recent developments in the field such as Dynamic Mode Decomposition (DMD) [81, 90, 91, 101, 107]. This latter aspect of our proposal is unique to our approach and endows the methodology with various attractive features such as the possibility of establishing certain bounds on uncertainty and retaining the underlying dynamic path of the dynamic phenomena.

Accordingly, we will utilize the temporal dynamics afforded by the temporal radiographic data at DARHT to constrain the solution to the manifold given by the hydrodynamic phenomena rather than simply using individual radiographic time slice to perform density reconstructions. Furthermore, we will utilize one of the great mathematical discoveries of the 20th century *i.e.*, Koopman Operator Theory, to facilitate this process [70]. This will enable a physics-based constraint on the temporal evolution of the density field that may allow for a potential breakthrough in our ability to perform high fidelity density reconstructions. Importantly, our consideration of the underlying physics permits us to make generalization and uncertainty estimates which would otherwise not be

possible using traditional neural network approaches.

Our development of a methodology that uses a technique first proposed by Koopman to effectively map the non-linear behavior of a dynamic system onto a linear space draws upon recent advances in computing the Koopman modes using Dynamic Mode Decomposition (DMD) and associated modifications [81,90,91,101,107]. Using the Koopman linear operator representation of the underlying physics facilitates the inclusion of hydrodynamics in the iterative framework for density reconstruction, and permits the examination of the parameter space with high accuracy and thereby reduced computational time to investigate a multitude of physics models and their associated parameters for EOS, material strength and high explosives.

Before presenting our promising results using this remarkable theory we briefly remark that indeed Koopman operator theory has had tremendous success in representing high complex physics in many fields including fluid dynamics turbulence and reactive flows [81, 90]. As an illustration we present the results of Schmidt whose pioneering work led to a resurgence in the use of the Koopman operator theory as given by his development of the DMD algorithm to find Koopman modes. In this example Schmidt demonstrated that DMD could capture the numerical simulation of a flame based on a variable-density fuel jet. The numerical simulation and the DMD result is illustrated in Figure 5 [91] .

Indeed, we have had tremendous initial success using DMD to represent the Koopman operator for many hydrodynamic systems of interest using both analytical solutions as well as solutions obtained using computational fluid dynamic codes *i.e.*, CTH [80]. Figure 6 illustrates that DMD can represent the canonical Sedov explosion [11].

Furthermore, we have demonstrated that we can then use other DMD representations of hydrodynamic simulations to match the major features *i.e.*, the outer edge location as well as the shock location, and identify degenerate solutions, see Figure 7. In addition, we have shown that interpolation of hydrodynamic solutions is possible with very high accuracy thereby reducing the computational complexity of computing many solutions. Indeed, we have shown great process in the ability to represent hydrodynamic behavior as described by the coupled partial differential equations with a mapping obtained via DMD to represent the density field. Accordingly, we present a framework for a coupling of a radiographic and hydrodynamic code in Figure 8.

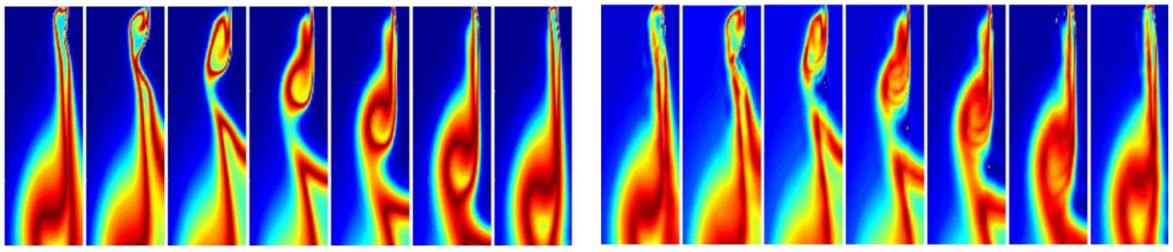

Figure 5: Selected (equispaced) snapshots of a passive tracer from a numerical simulation of a reactive flow. One shedding cycle is shown in a subregion of the full computational domain

3.4 Limited View Reconstruction Methods

Although reconstructions at DARHT have been performed for almost twenty years using the axisymmetric assumptions, almost every DARHT shot has exhibited clear evidence of three dimensionality. In fact, this was first recognized by Klasky where he pointed out this out to the BIE community [53, 55, 64]. Later Makaruk verified this finding [75–78].

While some work has been performed to circumvent the axisymmetric assumption used to perform reconstructions for the DAHRT data using the BIE, reconstructions still suffer due to this assumption [52,54,58,59,61,62]. The limited view radiographic data acquired at DARHT in conjunction with the axisymmetric assumption limit the application of traditional tomographic methods afforded by the use of filter back projection algorithms that require

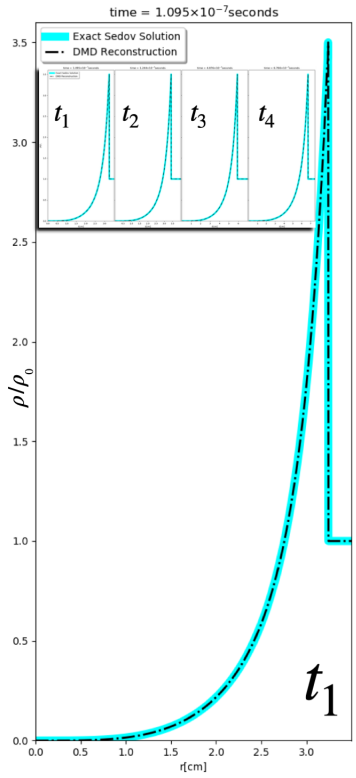

Figure 6: Density as a function of radius are shown for the exact Sedov solution and the DMD reconstruction in cyan and black, respectively for our set of four snapshots. The ambient density of $\rho_0 = 30.2$ was not in the set of solution space used to construct the initial DMD solution.

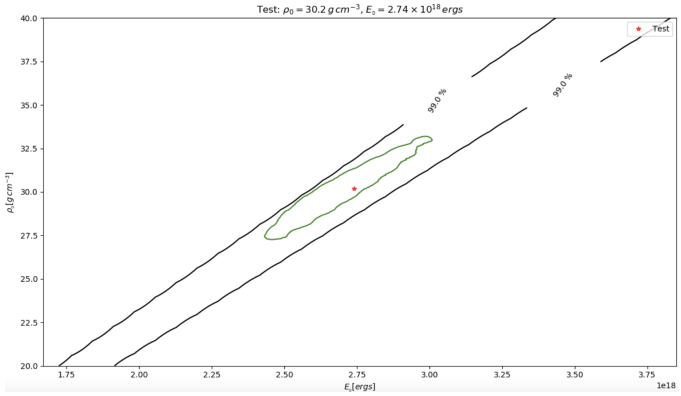

Figure 7: Confidence intervals are shown for density as a function of radius. The black curves denote the 99% confidence interval for the case when only the shock locations are used. The green contour shows the best fit for the 1% lowest χ^2 values for the case when the map is constructed using the DMD generator function with density increments $\Delta\rho_0 = 0.02$.

many views [31]. However, we can use *a priori* knowledge about the properties of the object to restrict the solution to the physically plausible domain. Indeed, we are currently investigating methods to enable the extraction of additional information from the data using priors as well as data-simulation fusion, for enabling additional information from the two-views of today, and potentially additional views of tomorrow, to allow for higher quality reconstructions.

We have begun this progression to more modern techniques to perform limited view reconstructions by ex-

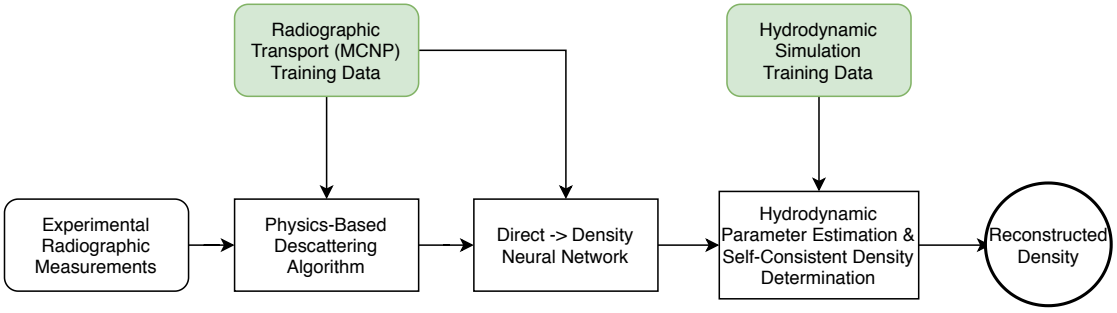

Figure 8: Flow-chart for EREBUS coupled Radiogrpahic and Hydrodynamic Reconstruction

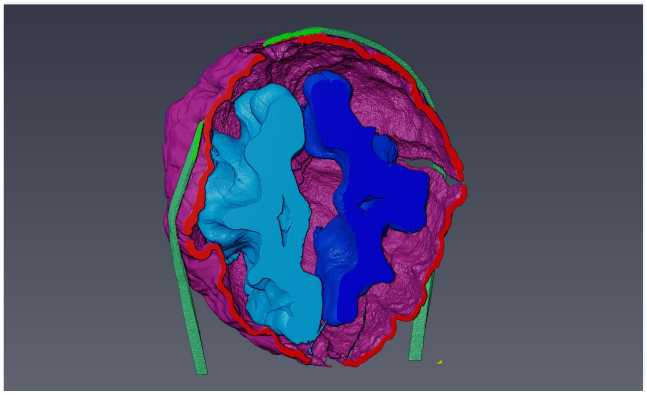

Figure 9: Tomographic Reconstruction of a walnut using FDK with segmentation.

amining a variety of canonical phantoms *i.e.*, Shepp Logan phantom [93], to X-cat phantom [10], as well as the reconstruction of walnuts using Micro-CT [27]. In Figure 9 we present a segmentation of tomographic reconstruction using a 4501 views using filtered back projection of a walnut using a standard FBP algorithm.

Traditional filtered backprojection reconstruction techniques have allowed for excellent reconstructions using a full complement of views, but this technique cannot, in general, provide for reconstructions with very limited data. Indeed, in Figure 10 we demonstrate the inability of the filtered backprojection technique to perform adequate reconstructions when there are limited views.

Indeed, this was a major issue when the number of views for the Advanced Hydrodynamic Facility was being evaluated. In fact, the assessment based on existing techniques *i.e.*, FDK eventually led to the abandonment of the concept and eventual refocusing to DARHT, [72]. However, methods to perform tomographic reconstructions with incomplete projections have proliferated in the last two decades. Starting from the Algebraic Reconstruction Technique (ART) [82] and later SART (Simultaneous Algebraic Reconstruction Technique) [3] as well as maximum likelihood expectation maximization algorithm have been proposed [103]. Other techniques including penalized likelihood reconstruction [33] and compressed sensing using total variation regularization total variation [94], as well as wavelet based techniques [74] to treat incomplete projections have allowed for improved reconstructions when limited projections are available [22, 46, 106].

Indeed, we have utilized a penalized least square method to perform reconstructions of the XCAT phantom as well as walnuts to achieve much better reconstructions than those achievable using standard filterd backprojection techniques [33]. These improved results for the walnut tomographic reconstructions may be observed from examination of in Figure 11

While there has been substantial improvement in the reconstructions using the traditional regularization approaches, a more recent trend is the application of deep learning (DL) to X-ray CT reconstruction [40, 108]. Cierniak combined a backprojection operation with a Hopfield neural network to reconstruct a CT image from projections [24]. Based on a persistent homology analysis, Han developed a deep learning residual architecture for sparse-view CT reconstruction [40]. Using multi-scale wavelets, they extended their work to limited angle CT

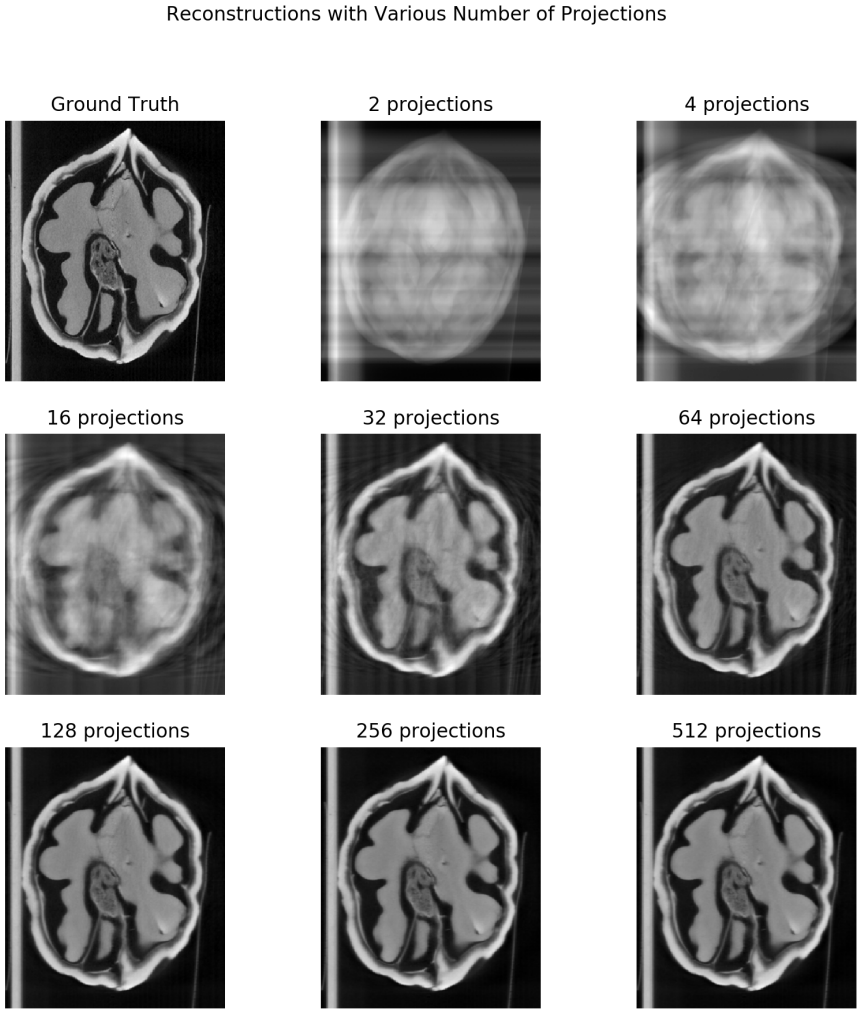

Figure 10: Filtered backprojection reconstruction of walnuts using various number of projections. Reconstruction fidelity increases with additional projections.

reconstructions. More recently, a new deep learning reconstruction framework for CT with incomplete projections utilized tight coupling of the deep learning U-net and FBP algorithm in the domain of the projection sinograms [29]. Indeed, DL methods described have shown tremendous potential in reducing the number of views and retaining excellent reconstruction ability [39, 48, 85]. Consequently, we have begun to employ these methods in addressing limited view tomographic problems. Indeed, Han of our team has utilized his residual learning network to demonstrate the tremendous improvements that can be afforded relative to traditional total variational approaches to limited view reconstruction, as shown in Figure 12 [40].

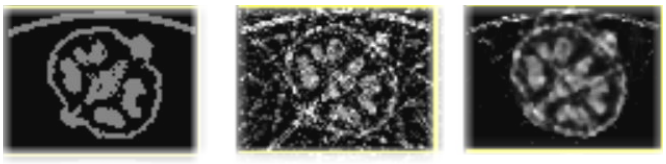

Figure 11: Limited-view tomographic reconstruction of walnut. (a) Ground truth of walnut slice. (b) FBP reconstruction from 16 views suffers from severe streaking artifacts. (c) Total-variation regularization mitigates streaking artifacts and yields a high-quality reconstruction.

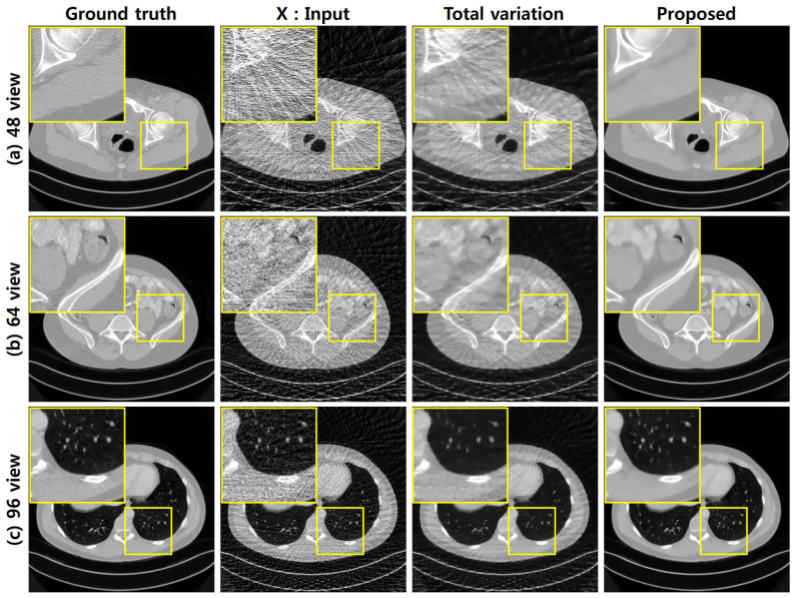

Figure 12: Limited-view tomographic reconstruction of the XCAT phantom. The first column shows three slices of the ground truth XCAT phantom. Traditional filtered backprojection (column 2) results in streaky artifacts due to the limited number of views. The second column shows a traditional filtered backprojection reconstruction. The third column shows the reconstruction obtained using total-variation regularization; while an improvement over FBP, streaky and patchy artifacts are present. The fourth column shows reconstruction using Deep Residual U-Net reconstruction [40]; these reconstructions contain significantly fewer artifacts than the total-variation reconstructions.

4 Integration with X-Division Code

We are currently working with the X-Div to utilize both software principles as well as interfaces to allow for integration of HART into RISTRA. The integration of these codes will more readily allow for the design community to better utilize the radiographic data as well as serve as a mechanism to improve the collaboration between the experimentalist/radiographic analyst, and the design community.

5 Uncertainty Analysis

We have recently proposed the use of our machine learning based reconstruction algorithm to address uncertainties in weapons performance in a principled scientific manner. Here we briefly we summarize the two proposed approaches.

While it is clear that a Bayesian framework is well suited to comprehensively quantify uncertainty in this context, the chief impediment to a practical realization of such an analysis is the immense computational cost of the physics-based forward model C_{fwd}. In the first approach, we will include the deep learning based forward model in a Bayesian inference framework [83] to establish uncertainties in the reconstruction. Again, recent improvements in computational inference methods will be brought to bear on this issue *e.g.*, by using Hamiltonian Monte Carlo whose computational cost only scales as $10^3 \times C_{fwd}$ as compared to Metropolis based schemes whose cost scales as $10^6 \times C_{fwd}$ [7].

In the second approach will use Bayesian machine learning and variational inference methodologies to improve upon the deep learning models. For example, variational inference methods have now been developed that brings the cost of probabilistic inference to levels comparable to that of point estimates [9]. Subsequent to this, we will utilize learned models for performance metrics to assess the propagation of these uncertainties to performance metrics.

6 X-Division Collaboration

With the curtailment of nuclear testing, the use of radiographic data is crucial in the certification of the stockpile as well as in other weapons related activities. Consequently, the full exploitation of the DAHRT radiographic data is crucial in improving our understanding of weapons physics. As such, it is imperative to obtain the highest quality of information from the radiographs as well as estimates in their uncertainty.

While the use of current radiograhic data in conjunction with the BIE has provided the opportunity to improve our qualitative understanding of hydrocode performance via qualitative density reconstructions, major physics improvements to the forward modeling as well as significant improvements in the priors *i.e.*, hydrodynamic priors, are necessary to enable the quantitative comparison of density fields from the radiographic data.

Currently, the radiographic data is used to verify high explosive (HE) models used in the codes. Most HE models parameters are based on experimental data from simple experiments. These parameters are sufficient for some geometries, but do not provide high levels of accuracy in complex geometries. It is therefore important to have radiographic reconstructions to verify shock locations for code predictions.

Radiographic data may also be utilized to examine modeling parameters for material strength as well as EOS models. However, the refinements in both the physics models utilized in the forward model as well as a more proper treatment in the temporal dynamics and asymmetries is necessary. These refinements to the radiographic reconstruction modeling will afford the designers the ability to make adjustments to the models based upon high fidelity reconstructions.

XTD-SS is currently lending support to provide 1D density fields to seed the machine learning process. The next steps will involve the generation of both two and three dimensional density fields. As the process evolves both models can be used to verify the density reconstruction process.

7 Software Engineering

Given the production community that this toolbox is targeting, we will follow good software engineering practices so that the toolbox will work readily with stakeholder communities. We will leverage our connections with the

ASC Ristra Next Generation Code Project and follow their standards for software quality assurance: Like Ristra, we will employ web-hosted Git repositories for source code revision control, git workflow with forks, branches, pull requests (PR) and code reviews, unit tests and automated regression testing, documentation of meetings using collaboration tools such as confluence.

Each team member will have a personal fork of the main project repository, and development of new features or bug fixes is generally done on a task-specific branch within that fork.

When changes for a new feature are considered ready, a pull request (PR) is submitted to the project's main repository. Creation of a PR automatically triggers a continuous integration (CI) system, which will run a set of builds and tests of the new code. Other team members are also notified of the new PR, and are given an opportunity to review and comment. The PR can be merged into the main repository only after the CI tests pass, and the reviewing team members have approved the merge.

The code review that takes place in the PR process helps to eliminate software defects at an early stage of the lifecycle. The CI testing ensures that the latest version of the code in the repository will always be tested and usable by other developers, both within this collaboration and for outside stakeholders. This workflow has been proven for a variety of application efforts developed under the Ristra Project and in collaborations outside Ristra. Indeed, this type of workflow is common in industry and academia, and is consistent with the software engineering practices being fielded across other ASC projects. As the proposed code will service these communities, it is natural for our team to partner with the CS experts within Ristra to adopt this type of software engineering environment.

8 Experience in Radiographic Analysis of Data

Combined, the team has over fifty years of experience in analyzing DARHT radiographic data as well as data taken at the U1A facility and FXR. During the past two decades, this team has analyzed over 90% of all of the DARHT shots. As such, the tremendous advantage that our team has in analyzing radiographic data is a major strength of the team and we will fully exploit this to our advantage as new techniques are developed and then tested by our experienced radiographic analysts. Furthermore, a key step in analyzing experimental data is the pre-processing. To this end we will utilize and improve upon the work by our team member Dr. Warthen. In addition, the experience in analyzing the data in conjunction with partnership with both the Engineering and J Divisions experimental facilities will enable real time feedback in proposed methods.

9 Collaborators

In developing an integrated framework for solving DARHT radiographic reconstruction problems, along with other radiographic data, we have assembled a world class team consisting of both Los Alamos National Laboratory staff, as well as researchers from outside the laboratory. Within the laboratory our team includes members from the theoretical design group XTD as well as collaborators from CCS to assist in the areas of computational hydrodynamics , theoretical hydrodynamics, and most importantly to address first hand the needs of the customer. In addition, our team has support from T-Division in the area of inverse methods, plug and play priors, as well as machine learning. Physics as well J Division provide both the experience in radiographic science, the forward modeling development, as well as beam physics. The Engineering Division brings both tomographic and radiographic experience as well as the ability to utilize the Microtron to support the development and validation of new physics models to bolster the forward model. Finally, the support from XCP provides the over code integration and software engineering support to develop a software package to support the laboratory.

In addition, we have extensive University collaborations to provide world class support as well as a pipeline of students that will be the future developers and users of the software. Finally, we also have had extensive conversations with our fellow researchers at the Lawrence Livermore, and AWE about their participation in supporting this effort.

10 Leveraged Activities

As discussed within this proposal to facilitate the development of HART we will leverage numerous currently funded activities. These activities are as follows:

- DTRA/NNSA Emergency Response Software
- LDRD Prioritizing the Prior
- J-Division Limited View Tomography
- J-5 Capability Funding for Beam Physics
- RISTRA Software Integration Support

11 FTEs and Budgets

To develop a comprehensive set of radiographic reconstruction tools for analysis of the DARHT radiographic data we plan to leverage numerous efforts that will contribute to the success of this Project. These efforts include the DTRA/NNSA effort to develop radiographic reconstructions for the Emergency Response Applications as well as efforts to develop limited view radiographic reconstruction applications funded by J-Division. In addition, we also have efforts to develop more fundamental efforts in algorithm development to address dynamic radiography and advanced priors for inverse problem solving. Finally, the software engineering and integration efforts will be supported by the ASC Program. These efforts will substantially reduce the cost to the hydro Program to develop a radiographic reconstruction code. Accordingly, our proposed budget and 3-5 year plan to support this effort is provided in Appendix B

12 Radiographic Analyst Requests

To ensure the seamless transition from the BIE to HART we have carefully considered the current radiographic user community, as over 90-95%, of the DARHT radiographic analyses have been performed by our team members. Here we highlight additional elements that may not have been addressed within the body of our proposal that are integral to the development of a production code as well as reference relevant sections of proposal.

- Our code will rely almost exclusively on well tested/peer reviewed algorithms from the computational imaging community. This will limit issues that have plagued the BIE for the past two decades. Furthermore, our team members are world leaders in the computational imaging community having developed many of the algorithms that are now commonly utilized in modern reconstruction algorithms including the GGMRF prior, plug and play prior, as well as physics models necessary for correct forward modeling.
- We will conform to reasonable software quality assurance standards, including testing. By partnering with XCP, RISTRA Project, we will adhere to modern software standards.
- Effective 3D functionality (i.e., all functionality available in 2D, runs within a reasonable amount of time, optimizable, etc.). As world leaders in computational imaging with experience in tomographic methods all ray tracing will be performed using validated codes. Furthermore,our extensive experience in modern parallelizable optimization techniques will ensure better performance than that achieved by the BIE. Furthermore, our demonstrated machine learning mappings will enable both more accurate reconstructions in orders of magnitude less time than traditional forward modeling.
- Direct translation of experimental parameters to tool inputs. As we have a multi-disciplinary team of computational imaging, physicists, and radiography/beam physicists having performed in excess of 95% of the radiographic reconstructions at DARHT as well as operated the accelerators at DARHT we will ensure the physical interpretability of our results as well as ties to physics input parameters.
- MPI (message passing interface); intrinsically parallelizable to use high performance computing (HPC) machines. Our codes will be developed with massively parallel architecture in mind. In addition our partnership with XCP (Ristra Project) will enable the use of software architecture to promote best software practices.
- Easy importing of external tools and libraries. Our computer archetecture will afford the ability to utilize external tools and libraries. Furthermore, the integration into RISTRA will enable unprecedented integration with X-Div software.

- Options for both graphical (GUI) and scripted interfaces. Both options for interfacing will be incorporated.
- Ability to integrate CAD models or other standard model file formats . Our team has developed routines to perform mappings of BIE ascii format files in transport codes as well as routines to utilize X-Division output dumps (Display). In addition, we have previously demonstrated the ability to take ABAQUS models as well as CAD models as inputs into our transport routines.
- Open source, industry accepted modern software language (avoidance of anything proprietary, or licensing fees, as much as possible) e.g. Python versus Matlab. Our team will not employ the software language of SmallTalk nor other Matlab or other propriety software. Furthermore, by partnering with XCP, RISTRA Project, we will adhere to modern software standards.
- Modular ability to read/write files supporting various hydro codes. We will continue to utilize Ensight and other processing codes in conjunction with our X-Division partners and code developers to enable the seamless use of X-Division code outputs. Furthermore, by partnering with the RISTRA Project we will also ensure integration of reconstructions into X-Division software for supplemental analyses.
- Ability to tie into or easily compare to X-hydrocodes. By partnering with XCP and working closely with XTD we plan to integrate HART into RISTRA thereby allowing for seamless integration and comparisons.
- Ability to be straightforwardly used on a variety of radiographic sources/platforms. We will utilize computer software that enables elements of the code to be run on both HPC as well as local networks.

A Appendix A Code Capability Summary

Table 1: Comparison of the capabilities in the different radiographic codes.

	BIE	BIE 2.0	HART
Forward Modeling Approach	X	X	X
Scatter Closure Relationship			X
Noise Treatment			X
Detector Model	Empirical	Empirical	X
Flat Field Treatment	Empirical	Empirical	Utilize FFT and Machine Learning
Spot Movement	Empirical	Empirical	Integrated Prior from Beam Diagnostics
Tomographic Capability			X
Plug and Play Priors			X
Hydrodynamic Prior			X
Modern Optimization Algorithms			X
Machine Learning Algorithms			X
MPI		X	X
Real Time Completion of Density Reconstruction			X
Modern Computing Language		X	X
AWE/LLNL Collaboration	X	X	X
University Collaboration			X
Other Agency Funding			X

B Appendix B Funding Request and Proposed Schedule of Work

We propose to develop the HART Toolbox using by leveraging our work. To develop our three major tools we propose a budget of 800 k/yr for 5 years. Figure 13 presents a Gantt chart to show the progress of the Project activities.

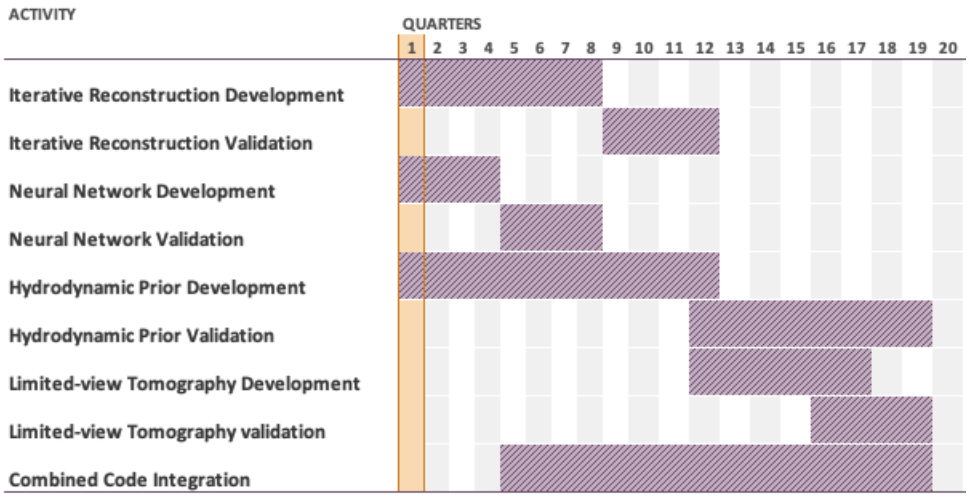

Figure 13: Gantt chart of the activities for HART development.

References

[1] Jan Aelterman, Hiêp Quang Luong, Bart Goossens, Aleksandra Pižurica, and Wilfried Philips. Augmented lagrangian based reconstruction of non-uniformly sub-nyquist sampled mri data. *Signal Processing*, 91(12):2731–2742, 2011.

[2] Sea Agostinelli, John Allison, K al Amako, John Apostolakis, H Araujo, P Arce, M Asai, D Axen, S Banerjee, G 2 Barrand, et al. Geant4—a simulation toolkit. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.

[3] Anders H Andersen and Avinash C Kak. Simultaneous algebraic reconstruction technique (sart): a superior implementation of the art algorithm. *Ultrasonic imaging*, 6(1):81–94, 1984.

[4] Inci M Baytas, Ming Yan, Anil K Jain, and Jiayu Zhou. Asynchronous multi-task learning. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pages 11–20. IEEE, 2016.

[5] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.

[6] M. Bertero. *Linear Inverse and Ill-Posed Problems*, pages 1–120. Advances in Electronics and Electron Physics Volume 75. Elsevier, 1989.

[7] Alexandros Beskos, Natesh Pillai, Gareth Roberts, Jesus-Maria Sanz-Serna, Andrew Stuart, et al. Optimal tuning of the hybrid monte carlo algorithm. *Bernoulli*, 19(5A):1501–1534, 2013.

[8] Christian H Bischof, Lucas Roh, and Andrew J Mauer-Oats. Adic: an extensible automatic differentiation tool for ansi-c. *Software: Practice and Experience*, 27(12):1427–1456, 1997.

[9] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.

[10] Jason Bond, Jack Frush, Sylvia Hon, Chris Eckersley, Cameron H Williams, Jianqiao Feng, Daniel J Tward, Tilak JT Ratnanather, MI Miller, D Frush, et al. Series of 4d adult xcat phantoms for imaging research and dosimetry. In *Medical Imaging 2012: Physics of Medical Imaging*, volume 8313, page 83130P. International Society for Optics and Photonics, 2012.

[11] DL Book. The sedov self-similar point blast solutions in nonuniform media. *Shock Waves*, 4(1):1–10, 1994.

[12] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011.

[13] A. Buades, B. Coll, and J. M. Morel. A non-local algorithm for image denoising. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 60–65 vol. 2, June 2005.

[14] Michael J Burns, George J Caporaso, Bruce E Carlsten, Yu-Jiuan Chen, Ken P Chow, Edward G Cook, Harold A Davis, Carl A Ekdahl, William M Fawley, Clifford M Fortgang, et al. Status of the dual axis radiographic hydrodynamics test (darht) facility. In *AIP Conference Proceedings*, volume 650, pages 139–142. American Institute of Physics, 2002.

[15] Richard H Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on scientific computing*, 16(5):1190–1208, 1995.

[16] Emmanuel Candès, Laurent Demanet, David Donoho, and Lexing Ying. Fast discrete curvelet transforms. *Multiscale Modeling & Simulation*, 5(3):861–899, 2006.

[17] Emmanuel J Candes and David L Donoho. New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities. *Commun. Pure Appl. Math.*, 57(2):219–266, 2004.

[18] Kathleen M Carley, Dave Columbus, and Ariel Azoulay. Automap user's guide 2012. Technical report, CARNEGIE-MELLON UNIV PITTSBURGH PA INST OF SOFTWARE RESEARCH INTERNAT, 2012.

[19] James L. Carroll. (u) preliminary radiographic uncertainty quantification for eos experiments at cygnus and fxr. Technical report, 2019. Sponsor: DOE.

[20] James L Carroll and Christopher D Tomkins. (u) towards quantifying uncertainties at darht. Technical report, 2013. Sponsor: DOE.

[21] James L Carroll and Christopher D Tomkins. (u) radiographic uncertainty quantification techniques for hydrotesting. Technical report, 2017. Sponsor: DOE.

[22] Guang-Hong Chen, Jie Tang, and Shuai Leng. Prior image constrained compressed sensing (piccs): a method to accurately reconstruct dynamic ct images from highly undersampled projection data sets. *Medical physics*, 35(2):660–663, 2008.

[23] Il Yong Chun, David Hong, Ben Adcock, and Jeffrey A. Fessler. Convolutional analysis operator learning: Dependence on training data. *CoRR*, 2019.

[24] Robert Cierniak. A new approach to image reconstruction from projections using a recurrent neural network. *International Journal of Applied Mathematics and Computer Science*, 18(2):147–157, 2008.

[25] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8):2080–2095, 2007.

[26] Viktor K Decyk. Upic: A framework for massively parallel particle-in-cell codes. *Computer Physics Communications*, 177(1-2):95–97, 2007.

[27] Henri Der Sarkissian, Felix Lucka, Maureen van Eijnatten, Giulia Colacicco, Sophia Bethany Coban, and Kees Joost Batenburg. A cone-beam x-ray ct data collection designed for machine learning. *arXiv preprint arXiv:1905.04787*, 2019.

[28] Jennifer Lynn Schei Disterhaupt and Marc Louis Klasky. Can blur effects be modeled by convolution in radiography. Technical report, 2018. Sponsor: USDOE National Nuclear Security Administration (NNSA).

[29] Jianbing Dong, Jian Fu, and Zhao He. A deep learning reconstruction framework for x-ray computed tomography with incomplete data. *PloS one*, 14(11), 2019.

[30] U Fano, LV Spencer, and MJ Berger. Penetration and diffusion of x rays. In *Neutrons and Related Gamma Ray Problems/Neutronen und Verwandte Gammastrahlprobleme*, pages 660–817. Springer, 1959.

[31] Lee A Feldkamp, Lloyd C Davis, and James W Kress. Practical cone-beam algorithm. *Josa a*, 1(6):612–619, 1984.

[32] Jeffrey A Fessler. Optimization methods for magnetic resonance image reconstruction: Key models and optimization algorithms. *IEEE Signal Processing Magazine*, 37(1):33–40, 2020.

[33] Jeffrey A Fessler and W Leslie Rogers. Spatial resolution properties of penalized-likelihood image reconstruction: space-invariant tomographs. *IEEE Transactions on Image processing*, 5(9):1346–1358, 1996.

[34] A.K. Fletcher, K. Ramchandran, and V.K. Goyal. *Wavelet denoising by recursive cycle spinning*, pages II–873–II–876. Institute of Electrical and Electronics Engineers, 2002.

[35] Roger Fletcher. *Practical methods of optimization*. John Wiley & Sons, 2013.

[36] Davis Gilton, Greg Ongie, and Rebecca Willett. Neumann networks for inverse problems in imaging. *CoRR*, 2019.

[37] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

[38] T Goorley, M James, Thomas Booth, F Brown, J Bull, LJ Cox, J Durkee, J Elson, Michael Fensin, RA Forster, et al. Initial mcnp6 release overview. *Nuclear Technology*, 180(3):298–315, 2012.

[39] Jawook Gu and Jong Chul Ye. Multi-scale wavelet domain residual learning for limited-angle ct reconstruction. *arXiv preprint arXiv:1703.01382*, 2017.

[40] Yo Seob Han, Jaejun Yoo, and Jong Chul Ye. Deep residual learning for compressed sensing ct reconstruction via persistent homology analysis. *arXiv preprint arXiv:1611.06391*, 2016.

[41] KM Hanson and GS Cunningham. The bayes inference engine. In *Maximum Entropy and Bayesian Methods*, pages 125–134. Springer, 1996.

[42] KM Hanson and GS Cunningham. Operation of the bayes inference engine. In *Maximum Entropy and Bayesian Methods Garching, Germany 1998*, pages 309–318. Springer, 1999.

[43] Laurent Hascoet and Valérie Pascual. The tapenade automatic differentiation tool: Principles, model, and specification. *ACM Transactions on Mathematical Software (TOMS)*, 39(3):1–43, 2013.

[44] Simon Hawe, Martin Kleinsteuber, and Klaus Diepold. Analysis operator learning and its application to image reconstruction. 22(6):2138–2150, June 2013.

[45] Luke Douglas Hovey, Jennifer Lynn Schei Disterhaupt, Marc Louis Klasky, and James Lamond Carroll. Darht radiographic spot blur study. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2019.

[46] Zhanli Hu, Juan Gao, Na Zhang, Yongfeng Yang, Xin Liu, Hairong Zheng, and Dong Liang. An improved statistical iterative algorithm for sparse-view and limited-angle ct image reconstruction. *Scientific reports*, 7(1):1–9, 2017.

[47] Kyong Hwan Jin, Michael T McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, 2017.

[48] Kyong Hwan Jin, Michael T McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, 2017.

[49] David W Juedes. A taxonomy of automatic differentiation tools. Technical report, Argonne National Lab., IL (United States), 1991.

[50] Donghwan Kim and Jeffrey A Fessler. Adaptive restart of the optimized gradient method for convex optimization. *Journal of Optimization Theory and Applications*, 178(1):240–263, 2018.

[51] Marc L Klasky. Examination of non-axi-symmetric density perturbations on density reconstructions (u). Technical report, 2010. Sponsor: DOE.

[52] Marc L Klasky. Analysis of hydro-shot 3648 (u). Technical report, 2011. Sponsor: DOE.

[53] Marc L Klasky. (u) asymmetries in hydroshots, implications on density reconstructions. Technical report, 2011. Sponsor: DOE.

[54] Marc L. Klasky. (u) comparisons of 3648 axis 1 and axis 2 time 3 density reconstructions. Technical report, 2011. Sponsor: DOE.

[55] Marc L Klasky. Evidence of three-dimensionality at darht and implications metric (u). Technical report, 2012. Sponsor: DOE.

[56] Marc L. Klasky. (u) an examination of the bucky grid on radiographic performance at darht. Technical report, 2018. Sponsor: DOE.

[57] Marc L. Klasky. (u) darht axis 2 physics based analysis and uncertainties. Technical report, 2018. Sponsor: DOE.

[58] Marc L. Klasky. (u) analysis of hydro-shot 3683/final forward model. Technical report, 2019. Sponsor: DOE.

[59] Marc L. Klasky. (u) analysis of hydro-shot 3685. Technical report, 2020. Sponsor: DOE.

[60] Marc L. Klasky, Jennifer Disterhaupt, Youzuo Lin, Dean L. Sanzo, Saiprasad Ravishankar, and Bert Iskender. Erebus physics based radiography. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2019.

[61] Marc L. Klasky and Jacob Mendez. (u) analysis of hydro-shot 4274/final forward model and shot report. Technical report, 2016. Sponsor: DOE.

[62] Marc L. Klasky and Jacob Mendez. (u) analysis of hydro-shot 4275/final forward model and shot report. Technical report, 2016. Sponsor: DOE.

[63] Marc L. Klasky, Trevor Wilcox, and Theodore Mockler. (u) 3683 comparisons of reconstructions and hydrodynamic calculations. Technical report, 2020. Sponsor: DOE.

[64] Marc Louis Klasky. (u) 3648 axis 1/2 comparisons and implications. Technical report, 2011. Sponsor: DOE.

[65] Marc Louis Klasky. A correct flat field model for darht. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2019.

[66] Marc Louis Klasky. A correct flat field model for darht. Technical report, 2019. Sponsor: USDOE National Nuclear Security Administration (NNSA).

[67] Marc Louis Klasky and Saiprasad Ravishankar. A new approach to solving radiographic inversion problems using a de-scattering algorithm. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2019.

[68] Marc Louis Klasky, Trevor Wilcox, Jennifer Lynn Schei Disterhaupt, Yoseob Han, Luke Anthony Pfister, Saiprasad Ravishankar, and Michael McCann. Towards a unified radiographic/hydrodynamic approach to obtaining accurate density from dynamic radiographic measurements. Technical report, 2020. Sponsor: USDOE National Nuclear Security Administration (NNSA), Defense Threat Reduction Agency (DTRA).

[69] Marc Louisand Klasky, Jennifer Lynn Schei, Youzuo Lin, Dean Lawrence Sanzo, Saiprasad Ravishankar, and Berk Iskender. Physics based machine learning for radiographic reconstructions. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2019.

[70] Bernard O Koopman. Hamiltonian systems and transformation in hilbert space. *Proceedings of the national academy of sciences of the united states of america*, 17(5):315, 1931.

[71] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

[72] P. W. Lisowski and J. A. Paisner. Ahf principles of design. *presented at the Proceedings of the 4th International Topical Meeting on Nuclear Applications of Accelerator Technology, Washington, D.C., 2001*.

[73] F. Luisier and T. Blu. Sure-let multichannel image denoising: Interscale orthonormal wavelet thresholding. 17(4):482–492, 2008.

[74] Xiaoqiang Luo, Wei Yu, and Chengxiang Wang. An image reconstruction method based on total variation and wavelet tight frame for limited-angle ct. *IEEE Access*, 6:1461–1470, 2017.

[75] Hanna E Makaruk. 3d perturbation of the inverse abel transform. Technical report, 2011. Sponsor: DOE.

[76] Hanna E Makaruk. (u) 3d effects in darht radiography: observations, analytical results and opportunities. Technical report, 2011. Sponsor: DOE.

[77] Hanna E Makaruk. Generalization of inverse abel transform 3d objects reconstruction from single radiographs for selected cases not fulfilling the axial symmetry assumption. Technical report, 2012. Sponsor: DOE.

[78] Hanna E Makaruk. Radiography: generalization of inverse abel transform 3d objects reconstruction. Technical report, 2012. Sponsor: DOE.

[79] Stéphane Mallat. *A Wavelet Tour of Signal Processing : The Sparse Way*. Academic Press, 2008.

[80] J Michael McGlaun, SL Thompson, and MG Elrick. Cth: a three-dimensional shock wave physics code. *International Journal of Impact Engineering*, 10(1-4):351–360, 1990.

[81] Igor Mezić. On applications of the spectral theory of the koopman operator in dynamical systems and control theory. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 7034–7041. IEEE, 2015.

[82] Klaus Mueller, Roni Yagel, and J Fredrick Cornhill. The weighted-distance scheme: a globally optimizing projection ordering method for art. *IEEE Transactions on Medical Imaging*, 16(2):223–230, 1997.

[83] Balasubramanya Nadiga, Chiyu Jiang, and Daniel Livescu. Leveraging bayesian analysis to improve accuracy of approximate models. *Journal of Computational Physics*, 394:280–297, 2019.

[84] Yurii Nesterov. A method for unconstrained convex minimization problem with the rate of convergence o (1/kˆ 2). In *Doklady an ussr*, volume 269, pages 543–547, 1983.

[85] Daniël M Pelt and James A Sethian. A mixed-scale dense convolutional neural network for image analysis. *Proceedings of the National Academy of Sciences*, 115(2):254–259, 2018.

[86] Zhimin Peng, Yangyang Xu, Ming Yan, and Wotao Yin. Arock: an algorithmic framework for asynchronous parallel coordinate updates. *SIAM Journal on Scientific Computing*, 38(5):A2851–A2879, 2016.

[87] Luke Pfister and Yoram Bresler. Learning filter bank sparsifying transforms. *IEEE Transactions on Signal Processing*, 67(2):504–519, 2019.

[88] Saiprasad Ravishankar and Yoram Bresler. MR image reconstruction from highly undersampled k-space data by dictionary learning. 30(5):1028–41, May 2011.

[89] DV Rose, DR Welch, BV Oliver, RE Clark, DL Johnson, JE Maenchen, PR Menge, CL Olson, and DC Rovang. Coupled particle-in-cell and monte carlo transport modeling of intense radiographic sources. *Journal of applied physics*, 91(5):3328–3335, 2002.

[90] Clarence W Rowley, Igor Mezić, Shervin Bagheri, Philipp Schlatter, and Dan S Henningson. Spectral analysis of nonlinear flows. *Journal of fluid mechanics*, 641:115–127, 2009.

[91] Peter J Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656:5–28, 2010.

[92] Martin E. Schulze and Michael Joseph Berninger. Energy spectrum and dose for axis ii. Technical report, 2018. Sponsor: USDOE National Nuclear Security Administration (NNSA).

[93] Lawrence A Shepp and Benjamin F Logan. The fourier reconstruction of a head section. *IEEE Transactions on nuclear science*, 21(3):21–43, 1974.

[94] Emil Y Sidky and Xiaochuan Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine & Biology*, 53(17):4777, 2008.

[95] Yu Sun, Brendt Wohlberg, and Ulugbek S Kamilov. An online plug-and-play algorithm for regularized image reconstruction. *IEEE Transactions on Computational Imaging*, 5(3):395–408, 2019.

[96] Hanlin Tang, Xiangru Lian, Ming Yan, Ce Zhang, and Ji Liu. D^2: Decentralized training over decentralized data. *arXiv preprint arXiv:1803.07068*, 2018.

[97] Afonso M Teodoro, José M Bioucas-Dias, and Mário AT Figueiredo. Image restoration and reconstruction using targeted plug-and-play priors. *IEEE Transactions on Computational Imaging*, 5(4):675–686, 2019.

[98] Christopher D. Tomkins. (u) 2018 review of the bayes inference engine (bie) and its application. Technical report, 2018. Sponsor: DOE.

[99] Christopher D Tomkins and James L Carroll. (u) radiographic data analysis with uncertainties using the bie. Technical report, 2017. Sponsor: DOE.

[100] Ivana Tosic and Pascal Frossard. Dictionary learning. 28(2):27–38, March 2011.

[101] Jonathan H Tu, Clarence W Rowley, Dirk M Luchtenburg, Steven L Brunton, and J Nathan Kutz. On dynamic mode decomposition: Theory and applications. *arXiv preprint arXiv:1312.0041*, 2013.

[102] Bert Vandeghinste, Bart Goossens, Roel Van Holen, Christian Vanhove, Aleksandra Pizurica, Stefaan Vandenberghe, and Steven Staelens. Iterative CT reconstruction using shearlet-based regularization. *Proc. SPIE 8313 Medical Imaging*, 8313:83133I–83133I–7, 2012.

[103] Yehuda Vardi, LA Shepp, and Linda Kaufman. A statistical model for positron emission tomography. *Journal of the American statistical Association*, 80(389):8–20, 1985.

[104] Singanallur V Venkatakrishnan, Charles A Bouman, and Brendt Wohlberg. Plug-and-play priors for model based reconstruction. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 945–948. IEEE, 2013.

[105] Bo Wahlberg, Stephen Boyd, Mariette Annergren, and Yang Wang. An admm algorithm for a class of total variation regularized estimation problems. *IFAC Proceedings Volumes*, 45(16):83–88, 2012.

[106] Ting Wang, Katsuhiro Nakamoto, Heye Zhang, and Huafeng Liu. Reweighted anisotropic total variation minimization for limited-angle ct reconstruction. *IEEE Transactions on Nuclear Science*, 64(10):2742–2760, 2017.

[107] Matthew O Williams, Ioannis G Kevrekidis, and Clarence W Rowley. A data–driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015.

[108] Tobias Würfl, Florin C Ghesu, Vincent Christlein, and Andreas Maier. Deep learning computed tomography. In *International conference on medical image computing and computer-assisted intervention*, pages 432–440. Springer, 2016.

[109] Qiong Xu, HY Yu, and XQ Mou. Low-dose x-ray CT reconstruction via dictionary learning. 31(9):1682–1697, September 2012.

[110] Mehrdad Yaghoobi, Sangnam Nam, Rémi Gribonval, Mike E Davies, et al. Analysis operator learning for overcomplete cosparse representations. In *European Signal Processing Conference (EUSIPCO'11)*, 2011.

[111] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, July 2017.

[112] G Zimmerman, D Kershaw, D Bailey, and J Harte. Lasnex code for inertial confinement fusion. Technical report, California Univ., Livermore (USA). Lawrence Livermore Lab., 1977. Classification Level (full-text): S.